

Process convolution approaches for modeling interacting trajectories

Henry R. Scharf

Department of Statistics, Colorado State University Fort Collins, CO, USA

Mevin B. Hooten

U. S. Geological Survey, Colorado Cooperative Fish and Wildlife Research Unit, Department of Fish, Wildlife, and Conservation Biology, Department of Statistics, Colorado State University, Fort Collins, CO, USA

Devin S. Johnson

Alaska Fisheries Science Center, Seattle, WA, USA

John W. Durban

Marine Mammal and Turtle Division, Southwest Fisheries Science Center, National Marine Fisheries Service, National Oceanic and Atmospheric Administration, La Jolla, CA, USA

Summary. Gaussian processes are a fundamental statistical tool used in a wide range of applications. In the spatio-temporal setting, several families of covariance functions exist to accommodate a wide variety of dependence structures arising in different applications. These parametric families can be restrictive and are insufficient in some situations. In contrast, process convolutions represent a flexible, interpretable approach to defining the covariance of a Gaussian process and have modest requirements to ensure validity. We introduce a generalization of the process convolution approach that employs multiple convolutions sequentially to form a “process convolution chain.” In our proposed multi-stage framework, complex dependencies that arise from a combination of different interacting mechanisms are decomposed into a series of interpretable kernel smoothers. We demonstrate an application of process convolution chains to model killer whale movement, in which the paths taken by multiple individuals are not independent, but reflect dynamic social interactions within the population. Our proposed model for dependent movement provides inference for the latent dynamic social structure in the study population. Additionally, by leveraging the positive dependence among individual paths, we achieve a reduction in uncertainty for the estimated locations of the whales, compared to a model that treats paths as independent.

Keywords: process convolution; kernel convolution; moving average; social network; animal movement

1. Introduction

Models for continuous random processes using kernel convolutions have come to be known in much of the spatial and spatio-temporal statistical literature as “process convolutions” (e.g., Higdon, 2002;

Address for correspondence: Henry Scharf, Department of Statistics, Colorado State University, 102 Statistics Building, Fort Collins, CO 80523, USA.

Email: henry.scharf@colostate.edu

Calder, 2007; Bolin and Lindgren, 2013). Also referred to as spatial moving averages (e.g., Cressie and Pavlicová, 2002), process convolutions first arose as a means for constructing valid covariance matrices for Gaussian processes (GPs), while relaxing the typical assumptions of stationarity and isotropy (e.g., Barry and Ver Hoef, 1996). One application area for process convolutions that has driven much of their recent theoretical development is for random processes taking place on stream networks (Ver Hoef et al., 2006). Stream networks occupy sub-manifolds of physical space on which specialized measures of distance often complicate enforcement of non-negative definiteness in covariance functions and have an inherent need for anisotropic dependence structures to accommodate the effects of direction in stream flow. Process convolutions offer an approach to modeling dependence that addresses both of these issues. We employ process convolutions to model the dependence observed in the trajectories of interacting animals.

A wealth of data arising from measurements of trajectories has spurred the development of many new statistical models for individual-based movement processes. These models have been used to study the individual movement patterns of a wide variety of animals, where scientific questions of interest are often focused on which exogenous environmental factors drive movement (Potts et al., 2014). A fundamental commonality of research related to animal movement is a need for accurate and precise knowledge about the paths traversed by each individual, leading to the development of many models used for reconstructing the true underlying movement processes from telemetry data (e.g., Fleming et al., 2015). Several excellent models for telemetry data have been proposed in both discrete- and continuous-time formulations that have sought to address the significant challenges making prediction of the true paths difficult (for a review see McClintock et al., 2014; Hooten et al., 2017). Telemetry devices are often subject to environmental conditions, and technological limitations that restrict their ability to deliver precise measurements of location at dense temporal resolutions. Thus, many analyses of animal movement must account for both measurement error and irregular temporal observations when estimating the true path traversed by an individual. In some cases, it may be reasonable to assume that the data include independent Gaussian noise; however, for many telemetry devices, it is necessary to account for complex, non-Gaussian forms of measurement error (Brost et al., 2015; Buderman et al., 2016). Irregularity in the frequency of telemetry observations can complicate the implementation of discrete-time models (although see Scharf et al., 2016, for one solution using a multiple-imputation procedure) and potentially lead to large uncertainty during time intervals with few or no observations.

Ecologists are also interested in the extent to which individuals in a population move in direct response to each other. Particularly in the case of large mammals, complex and dynamic social connections within a population can play a critical role in the movement behavior of individuals (e.g., Williams and Lusseau, 2006; Goldenberg et al., 2014; Scharf et al., 2016). For example, two individuals with a strong social connection may exhibit similar movement patterns and visit the same locations, resulting in positive dependence between their respective paths. Inference about the social ties within a population, and how those ties change over time, can provide valuable information about the behavioral ecology of a population. However, there are few existing methods available that explicitly account for interactions among individuals (e.g., Haydon et al., 2008; Codling and Bode, 2014; Langrock et al., 2014; Russell et al., 2016; Scharf et al., 2016). Especially when direct observation of a species is infeasible, careful analysis of data gathered from telemetry devices can reveal useful information about a population’s social structure (Scharf et al., 2016).

We propose a joint model for the movement of multiple individuals that expands the suite of statistical methodology available for studying animal movement. Our approach allows researchers to study dependence among the paths of a population of individuals that arises as the result of social interactions. Also, by taking into account the dependence between individuals in a population, our approach has the potential to reduce both bias and uncertainty in reconstructing trajectories in the study population compared to models that treat individuals independently. We obtain improvements in path reconstruction by modeling the true movement processes of the population conditioned on an unobserved, dynamic social network. The inferred network provides a description of the social

ties within a population and how those ties change in time.

In what follows, we present a generalized approach to process convolutions in which a random process is decomposed into a sequence of one or more smoothing kernels that are convolved with a white noise process. Multiple stages of smoothing allow us to model dependence in time and between individuals separately, resulting in a multivariate Gaussian process that captures the combination of effects. It is important to note that, while the kernels responsible for inducing temporal and path-wise dependence are constructed separately, the resulting covariance function is not separable in the geostatistical sense.

We provide a detailed introduction to our generalized approach to process convolutions in Sections 2.1 and 2.2 and construct our joint model for movement in Sections 2.3 and 2.4. After discussing the implementation details in Section 3, we demonstrate path reconstruction in a simulation study in Section 4. We apply our model to the study of telemetry data arising from the movement of killer whales in Section 5 and close with a summary of our findings and future directions.

2. Methods

2.1. A multiple-kernel convolution framework

We first outline a new flexible framework for the development of Gaussian process (GP) models based on a synthesis of ideas from geostatistics, multivariate time series, and trajectory modeling. Our hierarchical framework relies on the kernel convolution approach for modeling random processes, known, in the spatial and spatio-temporal statistical literature, as “process convolutions” (e.g., Higdon, 2002; Calder, 2007; Bolin and Lindgren, 2013). A mean zero random process $\mu(\cdot)$ is called a process convolution if it is constructed by convolving a continuous random process, $dB(\cdot)$, with a kernel function, h , over a domain, \mathcal{T} , so that

$$\mu(t) = \int_{\mathcal{T}} h(t, \tau) dB(\tau). \quad (1)$$

If the random process $dB(\cdot)$ is Gaussian, the resulting process $\mu(\cdot)$ will also be Gaussian with covariance function

$$\text{Cov}(\mu(t), \mu(t^*)) = \int_{\mathcal{T}} \int_{\mathcal{T}} h(t, \tau) \text{Cov}(dB(\tau), dB(\tau^*)) h(t^*, \tau^*) \quad (2)$$

(Higdon, 2002). It is common to define the process $dB(\cdot)$ to be Gaussian white noise, which yields the simplified covariance function

$$C(t, t^*) \equiv \text{Cov}(\mu(t), \mu(t^*)) = \int_{\mathcal{T}} h(t, \tau) h(t^*, \tau) d\tau \quad (3)$$

(although see Nychka et al. (2015) for an example of a spatial process convolution where dB has covariance specified through a Gaussian Markov random field). A process convolution with kernel function $h(\cdot, \cdot)$, as in (1), represents a smoothing of the process dB , and the kernel is therefore often referred to as a smoother. In what follows, we refer to h as both a kernel function and a smoother, interchangeably.

The function defined in (3) is guaranteed to be non-negative definite, and therefore a valid covariance function, if $\int_{\mathcal{T}} h(t, \tau) d\tau < \infty$ and $\int_{\mathcal{T}} h^2(t, \tau) d\tau < \infty$ for all t (Higdon, 2002). In many applications, it may be easier to specify the proper form of dependence in a GP through the form of a kernel smoother rather than that of a covariance function. For example, by constructing process convolutions with asymmetrical kernels, Ver Hoef et al. (2006) accounted for directional dependence in spatial processes arising on stream networks, where directional flow plays a crucial role. For certain choices of h , it is possible to express the covariance function analytically (e.g., Higdon, 2002; Paciorek, 2003; Ver Hoef and Peterson, 2010), although this is not necessary for non-negative definiteness. In the cases where an analytic solution to (3) is not available, one can evaluate

the integral numerically. The class of covariance functions constructed using kernel convolutions is general, containing many of the parametric families commonly used in geostatistical settings (e.g., exponential, Gaussian, spherical), as well as more flexible dependence structures. For example, Higdon (1998) and Paciorek (2003) showed that process convolutions allow researchers to model both anisotropic and non-stationary dependence by letting $h(t, \tau)$ vary with t . In principle, any covariance function evaluated over a discretized domain has a process convolution representation on the same discretized grid, where the kernel function's values over the grid may be defined using a decomposition of the covariance matrix (e.g., Cholesky).

Hooten and Johnson (2016) used process convolutions to develop new models for trajectories by convolving a Wiener process, rather than white noise, with different kernel functions. Using Wiener process convolutions, they were able to create realistic models for animal movement without the characteristic “roughness” found in Brownian motion and define a framework that incorporated other existing models for movement, such as that in Johnson et al. (2008). Brownian motion itself can be thought of as a process convolution where the kernel is a step function given by $h^{(bm)}(t, \tau) \equiv \mathbf{1}_{\{\tau \leq t\}}$ and the process is Gaussian white noise. Thus, the Hooten and Johnson (2016) framework can be viewed as a two-stage sequential smoothing procedure, where the position at time t is

$$\mu(t) = \int_{\mathcal{T}} h(t, \tau) \left(\int_{\mathcal{T}} h^{(bm)}(\tau, \tilde{\tau}) dB(\tilde{\tau}) \right) d\tau. \quad (4)$$

To account for Brownian motion with an initial position far from the origin, one may define the process $dB(\cdot)$ to be a continuous white noise process for all $t > 0$, with arbitrary initial variance so that $dB(0) \sim \mathcal{N}(0, \sigma_0^2)$. Hooten and Johnson (2016) provided several examples of possible kernel functions and demonstrated a computationally efficient procedure to incorporate temporal non-stationarity into the kernel, allowing them to fit highly flexible models to telemetry data arising from animal movement trajectories.

It is possible to write the two-stage process convolution in (4) in terms of a single effective smoothing kernel $\tilde{h}(t, \tau) = \int_{\mathcal{T}} h(t, \tau) h^{(bm)}(\tau, \tilde{\tau}) d\tau$ using Fubini's theorem to change the order of integration. Moreover, the equivalent, single-stage representation for $\mu(\cdot)$ is not limited to two-stage smoothing processes. One may specify any number of kernel functions and the resulting multi-stage kernel convolution can be written in terms of a single effective kernel. This nested structure motivates a novel model-building approach in which an arbitrarily large collection of convolution kernels is used to specify a valid GP. We refer to a stochastic processes resulting from the iterative convolution of a chain of kernels $h^{(1)}, \dots, h^{(L)}$ as a “process convolution chain” (PCC) and write it in its expanded form as

$$\begin{aligned} \mu^{(1)}(\tau_1) &= \int_{\mathcal{T}} h^{(1)}(\tau_1, \tau_0) dB(\tau_0) \\ \mu^{(\underline{2})}(\tau_2) &= \int_{\mathcal{T}} h^{(2)}(\tau_2, \tau_1) \mu^{(1)}(\tau_1) d\tau_1 \\ &\vdots \\ \mu^{(\underline{L})}(\tau_L) &= \int_{\mathcal{T}} h^{(L)}(\tau_L, \tau_{L-1}) \mu^{(\underline{L-1})}(\tau_{L-1}) d\tau_{L-1}. \end{aligned} \quad (5)$$

The underline in stages $2 \leq l \leq L$ (i.e., $\mu^{(\underline{l})}(\tau_l)$) distinguishes the PCC composed of all kernels $\{h^{(k)} : k \leq l\}$ from that of convolving $h^{(l)}$ with the process $dB(\cdot)$ directly (written $\mu^{(l)}(\tau_l)$). We use matching superscripts to denote the covariance function for a given process (e.g., $\text{Cov}(\mu^{(\underline{L})}(t), \mu^{(\underline{L})}(t^*)) = C^{(\underline{L})}(t, t^*)$). model A collapsed version of a PCC using a single effective kernel and the white noise process $dB(\cdot)$ can be written as

$$\mu^{(\underline{L})}(\tau_L) = \int_{\mathcal{T}} h^{(\underline{L})}(\tau_L, \tau_0) dB(\tau_0) \quad (6)$$

if we define the effective kernel as

$$h^{(\underline{L})}(\tau_l, \tau_0) \equiv \int_{\mathcal{T}} \cdots \int_{\mathcal{T}} h^{(l)}(\tau_l, \tau_{l-1}) \cdots h^{(1)}(\tau_1, \tau_0) d\tau_1 \cdots d\tau_{l-1}. \quad (7)$$

To simplify notation, we write $\tilde{h} = h^{(\underline{L})}$ for the effective kernel constructed from the entire chain, $\tilde{\mu}$ for the full process $\mu^{(\underline{L})}$, and \tilde{C} for the corresponding covariance function. In general, the kernels $h^{(k)}$ are not commutative, in the sense that

$$\int_{\mathcal{T}} \int_{\mathcal{T}} h^{(l_1)}(t, \tau) h^{(l_2)}(\tau, \tau_0) dB(\tau_0) \neq \int_{\mathcal{T}} \int_{\mathcal{T}} h^{(l_2)}(t, \tau) h^{(l_1)}(\tau, \tau_0) dB(\tau_0). \quad (8)$$

Thus, both the forms of the kernels and their order are important when specifying a PCC.

Specifying a GP through an ordered chain of smoothers allows for considerable flexibility. For example, if we specify a chain of size $L = 1$ with kernel function $h^{(1)}(t, \tau_1) \equiv \mathbf{1}_{\{\tau_1 < t\}}$, we recover simple Brownian motion. If we increase the length of the chain by including a second kernel, $h^{(2)}(t, \tau_2) = (1 + \frac{\tau_2 - t}{\phi}) \mathbf{1}_{\{-\phi < \tau_2 - t \leq 0\}}$, we recover a structure used by Ver Hoef and Peterson (2010) for modeling dependence in stream networks (also mentioned in Hooten and Johnson, 2016, Figure 2, row 3). Another specific case of PCCs are so-called $(k - 1)$ -fold integrated Wiener processes (Shepp, 1966; Wecker and Ansley, 1983; Rue and Held, 2005), which use $L = k - 1$ kernels of the form $h^{(l)}(t, \tau) = (t - \tau)/(l - 1)$ and allow one to model GPs with exactly m continuous derivatives. The multi-stage decomposition is similar in spirit to the way random processes with multiple dependence scales are decomposed additively in multi-resolution processes (Higdon, 2002; Nychka et al., 2015; Katzfuss, 2016); the difference is the convolution chain approach combines model components using convolutions, rather than additively.

In what follows, we construct convolution kernels that allow us to specify multivariate GPs; these provide a powerful tool for modeling multiple trajectories arising from interacting individuals. We return to this specific application of PCCs in Section 2.3.

2.2. Finite representation

The joint distribution of $\tilde{\mu} \equiv (\tilde{\mu}(t_1), \dots, \tilde{\mu}(t_n))'$ for $\mathbf{t} \equiv (t_1, \dots, t_n)' \subseteq \mathcal{T}$ is mean-zero Gaussian with covariance function $\tilde{C}(t, t^*) = \int_{\mathcal{T}} \tilde{h}(t, \tau) \tilde{h}(t^*, \tau) d\tau$ for any $t, t^* \in \mathbf{t}$ (if the integrals of \tilde{h} and \tilde{h}^2 exist). When this integral cannot be computed analytically, one can perform the integration numerically by selecting a fine grid of values $\boldsymbol{\tau} \equiv (\tau_0, \dots, \tau_m)'$ from the domain \mathcal{T} , and evaluating the discrete sum $\sum_{i=1}^m \tilde{h}(t, \tau_i) \tilde{h}(t^*, \tau_i) \Delta\tau_i$, where $\Delta\tau_i = \tau_i - \tau_{i-1}$. In the discretized setting, the process convolution may be written as a matrix product $\tilde{\mu} = \tilde{\mathbf{H}} \boldsymbol{\varepsilon}_0$, where the j^{th} row of $\tilde{\mathbf{H}}_{n \times m}$ is the function $h(t_j, \tau)$ evaluated at all $\tau \in \boldsymbol{\tau}$, and $\boldsymbol{\varepsilon}_0 \equiv (\varepsilon_0(\tau_1), \dots, \varepsilon_0(\tau_m))'$ with each $\varepsilon_0(\tau_i) \sim \mathcal{N}(0, \Delta\tau_i)$. Assuming the grid times $\boldsymbol{\tau}$ are equally spaced with intervals of size $\Delta\tau$, the joint distribution of $\tilde{\mu}$ can be expressed using matrix notation as

$$\tilde{\mu} \sim \mathcal{N}(\mathbf{0}, \Delta\tau \tilde{\mathbf{H}} \tilde{\mathbf{H}}'). \quad (9)$$

As the density of the grid grows to infinity ($m \rightarrow \infty$), the outer product in (9) approaches the covariance defined by the continuous integral over \mathcal{T} . Alternatively, one can follow the approach of Higdon (2002), who used finite process convolutions (also called discrete process convolutions in Calder, 2008) as an approximation, and choose $m < n$ to yield a fixed-rank model for the continuous process. Fixed-rank models can offer considerable computational efficiency and provide an adjustable level of implicit regularization (Wikle, 2010).

2.3. A process convolution chain for dependent movement

2.3.1. Social smoothing

We construct a novel model for animal movement using the PCC approach within a Bayesian hierarchical modeling framework. To describe the basic procedure, we first consider paths in one dimension. The model we describe can be readily extended to movement in two or more dimensions. We use a three-stage ($L = 3$) PCC that includes a specialized kernel function constructed to induce dependence among the paths of p different individuals. The inter-path dependence that arises is based on a weighted, undirected, latent social network (Goldenberg et al., 2010).

Let the random variable $w_{ij}(\tau) \in [0, 1]$ denote the connection weight between individuals i and j at time τ , with $w_{ii}(\tau) \equiv 1$ and $i, j \in \{1, \dots, p\}$. We propose a “social” smoothing kernel of the form

$$h_{ij}^{(soc)}(\tau_{soc}, \tau) \equiv \mathbf{1}_{\{\tau = \tau_{soc}\}} \frac{w_{ij}(\tau_{soc})}{|w_{i\cdot}(\tau_{soc})|}, \quad (10)$$

$$|w_{i\cdot}(\tau_{soc})| \equiv \sum_{j=1}^p |w_{ij}(\tau_{soc})|. \quad (11)$$

We write the kernel in (10) with explicit dependence on arbitrary time τ for completeness, but will hereafter omit the superfluous second argument and write $h_{ij}^{(soc)}(\cdot)$ for brevity. To provide intuition for the effect of convolving with the social smoothing kernel, we first consider the case of a multivariate process with two individuals. We assume that, for the process $\boldsymbol{\mu}(\cdot) \equiv (\mu_1(\cdot), \mu_2(\cdot))'$, the variables $\mu_1(\tau)$ and $\mu_2(\tau^*)$ are independent *a priori* for all τ and τ^* . After social smoothing, the resulting process for individual 1 is defined by

$$\mu_1^{(soc)}(\tau_{soc}) \equiv \sum_{j=1}^2 \int_{\mathcal{T}} h_{1j}^{(soc)}(\tau_{soc}, \tau) \mu_j(\tau) d\tau \quad (12)$$

$$= h_{11}(\tau_{soc}) \mu_1(\tau_{soc}) + h_{12}(\tau_{soc}) \mu_2(\tau_{soc}) \quad (13)$$

$$= \frac{\mu_1(\tau_{soc})}{1 + |w_{12}(\tau_{soc})|} + \frac{w_{12}(\tau_{soc}) \mu_2(\tau_{soc})}{1 + |w_{12}(\tau_{soc})|}. \quad (14)$$

From (13) and (14), it is clear that $\mu_1^{(soc)}(\tau_{soc})$ is a weighted average of the independent variables $\mu_1(\tau_{soc})$ and $\mu_2(\tau_{soc})$, where the weights are equal to $h_{11}(\tau_{soc})$ and $h_{12}(\tau_{soc})$, respectively. When the network connection is strong (i.e., $w_{12}(\tau_{soc}) \approx 1$), the two weights $h_{11}(\tau_{soc})$ and $h_{12}(\tau_{soc})$ will be approximately equal, and the effect of the kernel is to “squeeze” the previously independent paths toward a scaled version of their mutual mean. When the network connection is weak ($w_{12}(\tau_{soc}) \approx 0$), the weights will be close to 1 and 0 respectively, and the processes will be unaffected by the kernel, retaining the *a priori* independence between individuals. These effects are visible in the simulated process shown in the top plot of Figure 1.

The social smoothing kernel we propose is motivated by an interest in describing a mechanistic, interpretable driver of dependent movement. Specifically, it is useful for modeling the paths of individuals who have a tendency to move toward, and alongside, other individuals with whom they share connections. This type of dependence is ubiquitous in animal movement, but is by no means the only meaningful type of interaction. As an alternative, consider the movement of two highly territorial animals with overlapping territories in which observed paths would appear to “avoid” each other. To model mutual avoidance, we require a social kernel that “spreads apart” two otherwise independent paths when a strong social connection is present. In what follows, we focus on the particular kernel defined in (16), but we emphasize that PCCs, with suitably constructed kernels, offer a flexible way to specify models for a wide range of possible behavioral mechanisms.

Equipped with the social kernel given in (16), we outline the first two stages of our full three-stage PCC model for the dependent movement of p individuals. We begin with Gaussian white noise

and smooth at the first stage convolving $h_i^{(bm)}(t, \tau) \equiv \mathbf{1}_{\{\tau \leq t\}}$ with $dB(\cdot)$ to generate p independent instances of Brownian motion, each with its own initial position. At the second stage, we apply the social smoothing kernel $h_{ij}^{(soc)}(\cdot)$. Smoothing with the social kernel returns weighted averages of the Brownian processes, where the weights are proportional to the latent social weights $w_{ij}(t)$. Thus, two individuals i and j for whom $w_{ij}(t)$ is close to 1 will tend to have smoothed locations $\mu_i^{(soc)}(t)$ and $\mu_j^{(soc)}(t)$ that are close together in space. A third and final stage in the PCC, introduced in the following section, ensures that the random process has the proper temporal smoothness required for modeling animal movement.

2.3.2. Inertial smoothing

Marginally, the individual paths $\mu_i^{(soc)}$ generated from the two-stage PCC constructed from $h^{(bm)}$ and $h^{(soc)}$ are each an instance of Brownian motion. However, Brownian motion is unsuitable for direct modeling of animal movement because the instantaneous velocity of a particle traversing a Brownian path is discontinuous, and therefore the acceleration is not well defined. Discontinuities in the first derivative of the path processes imply that individuals are capable of instantaneous changes in their velocities and is inconsistent with the physical laws governing the mechanics of massive bodies (e.g., Feynman et al., 1963).

Several process convolutions have been explored that impart specific smoothness properties on a GP. For example, Shepp (1966), Wecker and Ansley (1983), and Rue and Held (2005) discuss applications of $(k - 1)$ -fold integrated Wiener processes, which ensure the existence of $k - 1$ continuous derivatives. Johnson et al. (2008) constructed a model for the movement of individual harbor and northern fur seals by specifying an Ornstein-Uhlenbeck process for the velocity, rather than the position. Hooten and Johnson (2016) modeled the true locations of an individual animal using a two-stage PCC in which the second kernel function is Gaussian, and Buderman et al. (2016) use a similar model with cubic splines as their kernels. In all of these examples, smoothness is imparted to a process through convolution with a kernel either implicitly or explicitly. We introduce a final stage of smoothing to yield paths that are guaranteed to be continuously differentiable.

Enforcing the existence of a continuous derivative for the movement processes ensures that individuals moving through the paths μ have an interpretable acceleration at every time t . Therefore, we refer to the final kernel function as an “inertial” smoother because its function is to generate paths with the physical properties required to obey the laws of classical mechanics. Following Hooten and Johnson (2016), we specify a Gaussian kernel with an unknown variance parameter, ϕ_{inl}^2 , as the inertial smoother. The effect of this final kernel, $h^{(inl)}$, is visible in the bottom plot of Figure 1.

2.3.3. The full three-stage process convolution chain

The expanded specification of the full three-stage PCC is given by

$$\mu_i^{(bm)}(\tau_{bm}) = \int_{\mathcal{T}} h_i^{(bm)}(\tau_{bm}, \tau_0) dB_i(\tau_0) \quad (15)$$

$$\mu_i^{(soc)}(\tau_{soc}) = \sum_{j=1}^p h_{ij}^{(soc)}(\tau_{soc}) \mu_j^{(bm)}(\tau_{soc}) \quad (16)$$

$$\tilde{\mu}_i(\tau_{inl}) = \mu_i^{(inl)}(\tau_{inl}) = \int_{\mathcal{T}} h^{(inl)}(\tau_{inl}, \tau_{soc}) \mu_i^{(soc)}(\tau_{soc}) d\tau_{soc} \quad (17)$$

where

$$h^{(inl)}(\tau_{inl}, \tau_{soc}) \equiv \exp\left(\frac{-(\tau_{inl} - \tau_{soc})^2}{2\phi_{inl}^2}\right). \quad (18)$$




Fig. 1. An example of a one dimensional GP arising from a three-stage social PCC for two individuals (light and dark). All three intermediate stages of smoothing are shown, beginning with Brownian motion (top, dotted), followed by “social” smoothing (top and bottom, dashed), and finally “inertial” smoothing (bottom, solid). The thick gray line in the top plot is proportional to the strength of the social tie $w_{12}(t)$ at all times t , where the extent of the y-axis corresponds to a range of $[0, 1]$. The thick gray line in the bottom plot is proportional to the inertial smoothing kernel $h^{(inl)}(t)$.

One of the primary benefits of using a PCC framework is that it allows one to decompose complex dependencies into an iterative sequence of relatively simple mechanisms. Positions of multiple individuals can, in general, exhibit dependencies in both time and among individuals, and the characteristics of these dependencies may themselves be dynamic. We deliberately construct a PCC using kernels that compartmentalize the mechanisms generating the complex dependencies in the collective movement of interacting individuals. The first stage of smoothing is in time, and not across individuals, so that $\mu_i^{(bm)}(\tau_{bm})$ and $\mu_j^{(bm)}(\tau_{bm})$ are independent for $i \neq j$. The second stage of smoothing is across individuals, and not in time, in the sense that $h_{ij}^{(soc)}(t, t^*) = 0$ for $t \neq t^*$. The final stage of smoothing is in time only.

Restricting individual smoothing components to operate in only one “dimension” at a time (in this case, either temporal or social) is not necessary. However, a primary motivation for constructing GP models in the PCC framework is to decompose a complex mechanism into components that are easier to understand separately. Therefore, it is natural to form PCC models with kernel functions that operate in a limited number of dimensions simultaneously.

Let $s_i(t)$ denote the observed position of individual $i \in \{1, \dots, p\}$ at time $t \in \mathcal{T}$. We model the true, unobserved position of individual i at time t proportional to $\tilde{\mu}_i(t)$ and add a measurement error process, $\varepsilon_i(t)$, to yield the data model

$$s_i(t) \equiv \sigma_\mu \tilde{\mu}_i(t) + \varepsilon_i(t). \quad (19)$$

The parameter σ_μ scales the random process $\tilde{\mu}$ to account for the overall spatial domain of the locations. We model the measurement error $\varepsilon_i(t)$ as i.i.d. mean-zero Gaussian random variables with variance σ_s^2 , although more complex measurement error processes could also be used (e.g., Brost et al., 2015; Buderman et al., 2016).

2.4. Process prior for the dynamic social network

The general model for dependent movement is highly parameterized. To evaluate the covariance $\Delta\tau \tilde{\mathbf{H}}\tilde{\mathbf{H}}'$ for a grid of m times, one must estimate $\binom{p}{2}m$ parameters associated with the social network that drives the second stage of the PCC. As we describe in Section 3, we fit the proposed model to data within a Markov chain Monte Carlo (MCMC) paradigm, which requires repeated inversion of a large, dense covariance matrix. Without further structure for $w_{ij}(\cdot)$, estimating the underlying network would be computationally infeasible. We increase computational efficiency by leveraging reasonable assumptions about the underlying network to constrain the space of possible latent social networks, thereby reducing the effective number of parameters in the model. We impose constraints on the network’s marginal complexity at each fixed time as well as the smoothness of its temporal evolution.

Our approach makes use of recently developed methods for modeling dynamic stochastic networks. To reduce the dimensionality of the latent network at a fixed time, we follow the approaches of Hoff et al. (2002), Hoff (2008), and Durante et al. (2014), who modeled social connections using a latent space. Latent space models for stochastic networks proceed by defining network connections through functions of the locations of particles in a latent d -dimensional “social” space. As an example, consider the following latent space model for the edges in a weighted, undirected, dynamic social network. Suppose the latent variables $\mu_i^{(w)}(t)$, $\mu_j^{(w)}(t)$, and $\mu_k^{(w)}(t)$ are positions in \mathbb{R}^2 , and edge weights are defined by $w_{ij}(t) = e^{-d_{ij}(t)}$, where $d_{ij}(t)$ is the Euclidean distance between $\mu_i^{(w)}(t)$ and $\mu_j^{(w)}(t)$. Defining edge weights based on the distance between particles in a latent space induces positive dependence among the connections $w_{ij}(t)$, because the latent distances are constrained by the characteristics of the latent space and the measure of distance. That is, if $d_{ij}(t)$ and $d_{jk}(t)$ are small, resulting in edge weights $w_{ij}(t)$ and $w_{jk}(t)$ that are close to 1, the triangle inequality implies $d_{ik}(t)$ must also be small, and $w_{ik}(t)$ will also be close to 1. Therefore, latent space approaches for modeling social networks can be understood as a way to induce a tendency to complete triangles

in a network. This phenomenon is often observed in human social networks (e.g., friends of friends tend to themselves be friends) (Hansell, 1984), and is also reasonable for many applications to animal social networks. A latent space approach to constraining the underlying social network has the computational advantage of reducing the number of parameters to estimate by a factor of p , because we only need to estimate p latent paths instead of $\binom{p}{2}$ pairwise relationships.

We define the latent social connection between individuals i and j at time τ_{soc} using the latent positions and an appropriately chosen functional g by

$$w_{ij}(\tau_{soc}) \equiv g\left(\boldsymbol{\mu}_i^{(w)}(\tau_{soc}), \boldsymbol{\mu}_j^{(w)}(\tau_{soc})\right). \quad (20)$$

In general, the type of network desired for the application (e.g., binary, weighted) may inform the particular choice for g . As discussed in Section 2.3, we construct the kernel $h_{ij}^{(soc)}$ under the assumption that the weights in the social network have support $[0, 1]$. The compact support of $w_{ij}(\cdot)$ motivates our choice of g , defined by $g(\mathbf{x}, \mathbf{y}) \equiv e^{-\|\mathbf{x}-\mathbf{y}\|_2^2}$. The functional g maps two vectors in the latent space to the unit interval $[0, 1]$ and closely follows the methods used for latent space network modeling by Hoff (2008) who showed that this construction induces positive dependence in the edges $w_{ij}(\cdot)$.

$$\boldsymbol{\mu}_i^{(w)}(\tau_w) \equiv \sigma_w \int_{\mathcal{T}} h^{(w)}(\tau_w, \tau) d\mathbf{B}_{w,i}(\tau), \quad (21)$$

$$h^{(w)}(\tau_w, \tau) \propto e^{\frac{-(\tau-\tau_w)^2}{\phi_w^2}}, \quad (22)$$

and $d\mathbf{B}_{w,i}(\cdot)$ are instances of two-dimensional Brownian motion, now used to define the positions in the latent social space. The parameter σ_w controls the dispersion of the latent paths $\boldsymbol{\mu}_i^{(w)}$, and is therefore related to the overall density of the social network, with smaller σ_w corresponding to higher connectivity. The parameter ϕ_w^2 controls the smoothness of the trajectories through the latent social space, and is related to the temporal stability of the network over time, with larger values of ϕ_w^2 corresponding to a more stable network. As described in Section 2.2, we approximate the continuous stochastic processes $\boldsymbol{\mu}_i^{(w)}(\cdot)$ using a set of independent normal random variables anchored at a finite number of knots yielding

$$\boldsymbol{\mu}_i^{(w)} \sim \mathcal{N}\left(\mathbf{0}, \sigma_w^2 \Delta t \tilde{\mathbf{H}}_w' \tilde{\mathbf{H}}_w\right). \quad (23)$$

3. Model implementation

We obtain realizations from the posterior distribution of the model parameters using a Markov chain Monte Carlo (MCMC) algorithm. In this section, we briefly discuss the most relevant features of model implementation. A more detailed description, including our choice of priors and hyperparameters, is available in Appendix 8.2.

We specify independent normal priors for the initial-location parameters and a conjugate hyperprior on the population-level variance so that $\mu_{0i} \sim \mathcal{N}(0, \sigma_0^2)$ and $\sigma_0^2 \sim \text{IG}(a_0, b_0)$. Additionally, we integrate the true location process $\tilde{\boldsymbol{\mu}}$ out of the likelihood, which allows us to avoid sampling the continuous movement process directly. This is a common technique used in spatial statistics (e.g., Gelfand et al., 2003; Finley et al., 2013) to improve mixing for the remaining parameters in the model. The underlying movement process $\tilde{\boldsymbol{\mu}}$ can be recovered *post hoc* using composition sampling (Finley et al., 2013). The resulting integrated model formulation is

$$\mathbf{s}_i(t) | \boldsymbol{\varepsilon}_w, \phi_{inl}, \sigma_s^2, \sigma_\mu^2, \sigma_0^2 \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}) \quad (24)$$

$$\boldsymbol{\Sigma} \equiv \sigma_s^2 \mathbf{I} + \sigma_\mu^2 \Delta \tau \tilde{\mathbf{H}}' \tilde{\mathbf{H}}, \quad (25)$$

where the parameters ϕ_{inl} , σ_0^2 , and ε_w enter the density through the definition of $\tilde{\mathbf{H}}$ in (15) - (18). The full-conditional distributions for the latent paths $\boldsymbol{\mu}_i^{(w)}$ are not conjugate. To update each latent-space path, $\boldsymbol{\mu}_i^{(w)}$, we employ a Metropolis-Hastings algorithm.

4. Simulation

As the dependence between individuals in a population weakens (i.e., $w_{ij}(t) \rightarrow 0$ for all $t, i \neq j$), our proposed model reduces to one that considers each path independently, but shares information about all non-social parameters across individuals. The PCC model constructed from a fixed, empty social network presents a natural baseline for comparison with our full model for dependent movement. Estimating the latent social network that gives rise to the dependence among paths represents the majority of the computational effort to fit our proposed model, and grows rapidly with the number of individuals under study. Thus, there is a natural incentive to model movement under the assumption of path-independence unless it can be shown to be deficient. We compare the performances of the full model for dependent movement (IP-DEP) with the special case of the model under the assumption of inter-path independence (IP-IND). In many cases, we find that accounting for the dependence among individuals results in significantly improved reconstruction of the true underlying paths.

Telemetry devices are subject to a wide range of environmental conditions that frequently result in large intervals of time during which no observed locations are recorded. Estimates of the true locations, $\tilde{\boldsymbol{\mu}}$, during long gaps between observations are often accompanied by large amounts of uncertainty that present a challenge for researchers studying animal behavior. In simulation, we show that the presence of moderate to strong dependence among observed paths provides an opportunity for improvement in path reconstruction when we take into account the joint distribution of all observed individuals. To illustrate the potential gains, we consider a simple setting of two individuals in which the connection status, w_{12} , is constant over the period of observation. We assume regular, uninterrupted observations for individual 1, but a gap in time, \mathcal{T}_g , exists in the observed sequence of telemetry locations for individual 2.

We evaluate the quality of a given model for path reconstruction in terms of both accuracy and precision. To assess the accuracy of a path reconstruction, we define a loss function termed the “squared path error” (SPE) that quantifies the agreement between a specific path reconstruction $\hat{\boldsymbol{\mu}}$ and the true path $\tilde{\boldsymbol{\mu}}_{true}$ as

$$\text{SPE}(\hat{\boldsymbol{\mu}}; \mathcal{T}_g, \tilde{\boldsymbol{\mu}}_{true}) = \frac{1}{|\mathcal{T}_g|} \int_{\mathcal{T}_g} \|\tilde{\boldsymbol{\mu}}_{true}(t_g) - \hat{\boldsymbol{\mu}}(t_g)\|_2^2 dt_g. \quad (26)$$

We define the reconstruction $\hat{\boldsymbol{\mu}}$ to be the posterior mean for $\tilde{\boldsymbol{\mu}}$ so that (26) represents a measure of how accurately the center of the posterior distribution matches the true underlying path. To assess the precision associated with a model, we compute the radii for circular 95% credible regions surrounding each point in the path $\hat{\boldsymbol{\mu}}(t_g)$ and average across the entire temporal gap in observations, \mathcal{T}_g , to yield an overall summary of precision we term the average circular 95% credible region radius (ACRR_{0.95}).

To provide a general sense of when accounting for path dependence results in the greatest gains in path reconstruction, we vary the strength of dependence, w_{12} , the proportion of the study interval made up by the observation gap, $|\mathcal{T}_g|/|\mathcal{T}|$, and the tortuosity of the paths, and use SPE in tandem with ACRR_{0.95} to compare the performance of IP-DEP and IP-IND. We vary tortuosity through the range parameter ϕ_{inl} that appears in the inertial smoother, with smaller values corresponding to more tortuous paths on average. For each combination of w_{12} , \mathcal{T}_g , and tortuosity, we used the five step procedure:

- (1) Simulate a realization of $\tilde{\boldsymbol{\mu}}_{true}$ and \mathbf{s} from our proposed model.
- (2) Fit both the dependent and independent-paths models to \mathbf{s} .

- (3) Sample 1000 paths from each of the two posterior distributions using composition sampling.
- (4) Compute the path error defined in (26) for the posterior mean.
- (5) Use the same 1000 draws from the posterior to compute the average radius of a circular 95% credible region surrounding a point $\mu_2(t_g)$.

We repeat the procedure 20 times for each combination of parameters to obtain an estimate of the variability between realizations of the simulated paths.

To facilitate direct comparison between the two models under consideration, we examine the ratios of SPE and $\text{ACRR}_{0.95}$ for IP-DEP and IP-IND, defining ratios such that the relevant value for IP-DEP appears in the denominator and values greater than 1 show support for IP-DEP. Figure 2 displays the ratios of SPE (top row) and $\text{ACRR}_{0.95}$ (bottom row) under all combinations of values for the parameters w_{12} , $|\mathcal{T}_g|/|\mathcal{T}|$, and ϕ_{inl} (tortuosity). Individual plots show the median ratio across all simulations for high (red, solid) and low (blue, dashed) tortuosities, with an associated polygon delineating the 25% and 75% quantiles across simulations. Columns organize the plots by gap size, and the strength of the social connection w_{12} increases along the x-axis within each individual plot.

Several general observations can be made about the circumstances under which fitting the full model for dependent movement offers the greatest improvements in path reconstruction. First, both precision and accuracy improve near monotonically with increasing w_{12} . Second, the greatest gains come when the gap size is moderate (column 3). As the size of the gap is continues to grow (column 4), stronger dependence between two paths is required to provide the same increases in accuracy and precision. Similarly, when the gap is brief, the difference between the reconstructed paths for each model is modest. Finally, the gains in performance for the full model are slightly greater for the case of high tortuosity.

5. Killer whales

We analyzed telemetry data for four killer whales near the Antarctic Peninsula (see Figure 3) over the course of five days in February 2014. Geographic positions were measured using Argos telemetry tags (for a complete description of the tags and study area see Andrews et al., 2008; Durban and Pitman, 2012). Although multiple types of killer whales have been described in this area, all four tags were deployed on individuals from the same population of the most common type of “B2” killer whales (Durban et al., 2016).

It is immediately apparent from the data that individuals 1 and 2 (Figure 3, bottom left) show potential evidence of a close connection. In addition, there is some ambiguity about the relationship between individuals 3 and 4 (Figure 3 top right), because they occupy approximately the same spatial region during the study period. In contrast, there is little reason to suspect dependence between the two pairs of individuals. We analyzed the movement of all four individuals jointly, which allowed for pooling of information about measurement error across the entire group, and provided an opportunity for basic model validation. By fitting a model for the joint movement of all four individuals with a fully flexible latent network structure, we were able to check for the presence of potentially spurious network connections, as we have a strong, *a priori*, belief that the underlying network should exclude connections between the two subgroups.

The top plot of Figure 4 shows the observation times for the killer whales in our study, with darker regions corresponding to a denser rate of telemetry measurements in time. Two different day-long gaps in observation occur for individual 1 (green, Figures 4 and 3), a consequence of the original study design used in the deployment of the telemetry tags which sought to balance the need for temporally dense observations with limitations on the battery life of the tags by collecting measurement only every other day for select individuals. The observation times for the other three individuals cover these gaps, suggesting that modeling the four paths jointly could allow for more precise and/or accurate estimates of the true path taken by individual 1 if moderate to strong dependence exists between it and any of the other three whales.

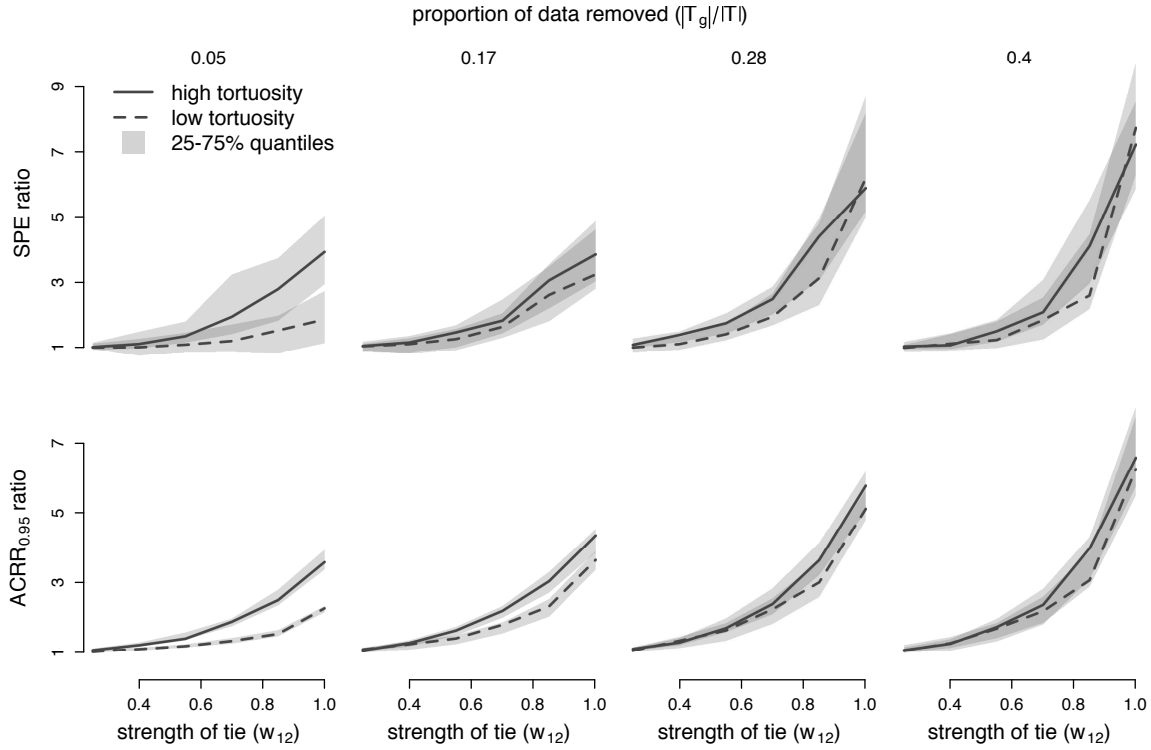


Fig. 2. Path errors. The top row of plots shows the ratios of SPE for the independent (IP-IND) to full (IP-DEP) models (thus, larger ratios correspond to the full model outperforming the alternative). The bottom row shows the analogous ratios of $\text{ACRR}_{0.95}$. Each column represents a fixed value for the proportion of the study interval made up by the gap in observations ($|\mathcal{T}_g|/|\mathcal{T}|$). Within each plot, the strength of the social connection increases along the x-axis. The lines represent the median value of all ratios across the 20 simulations, and the associated polygons represent the 25% and 75% quantile boundaries. Finally, the solid lines correspond to simulations with high tortuosity in the true paths, while dashed lines correspond to simulations with low tortuosity.

As in the simulation study (Section 4), we investigated path reconstructions generated by our proposed full model for dependent movement (IP-DEP), as well as the analogous model under an assumption of inter-path independence (IP-IND). We found that the full model generates path reconstructions in which the uncertainty about the position of individual 1 is dramatically reduced, compared to the reconstructions generated under the assumption of inter-path independence as measured by 95% circular credible interval radii (bottom plot of Figure 4). Circular credible regions for the true position of the whale during the two large gaps in observation occurring on February 12th and 14th for the IP-IND model have large radii, sometimes exceeding 70km. In contrast, circular credible regions for the true position based on the IP-DEP model are less than a third this size, and on par with the uncertainties for individuals for which we have dense observations. Despite the uncertainty about the latent social network, the IP-DEP model offers a substantial reduction in uncertainty about the true path taken by individual 1.

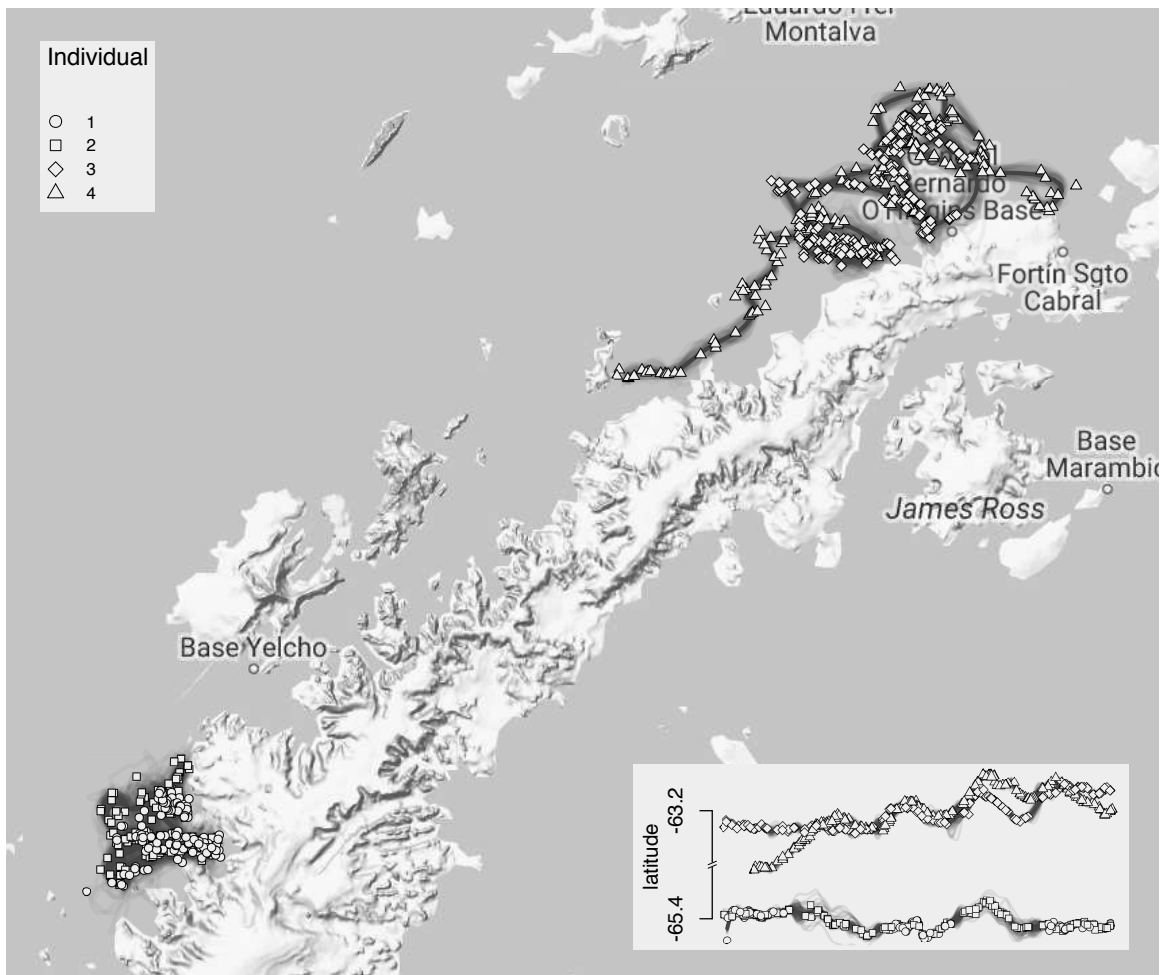


Fig. 3. Observed telemetry data and joint posterior distribution of the true paths of all individuals ($\tilde{\mu}$). The solid lines represent the posterior mean, and the semi-transparent lines are draws from the posterior distribution to illustrate uncertainty in the path reconstruction. The points are the observed locations from the Argos satellite system. The shapes correspond to the four individuals in the study, and match those used in Figure 4. The subplot shows the latitude of each individual over time. Map created with Kahle and Wickham (2013). Map data ©2017 Google.

The posterior distributions for the two network relationships of primary interest are shown in Figure 6. We can see strong evidence for a significant relationship between individuals 1 and 2, however there is essentially no evidence of a meaningful social connection between individuals 3 and 4. Similarly, there was essentially no evidence of connections existing between any other pairs of individuals (see Appendix 4 for all other relationships). There is large uncertainty about the strength of the connection between individuals 1 and 2, although the 95% posterior credible interval suggests that the two individuals are at least weakly interacting for the majority of the study period. The amount of uncertainty surrounding this connection is unsurprising, given the day-long gaps in observations for individual 1.

6. Discussion

Appropriate models for dependent, multivariate data are application-dependent. Gaussian processes offer a flexible, parsimonious tool for analyzing complex data, such as those that arise from measuring the movement of animals, if one can specify a valid form for the covariance. The need for covariances that are mathematically sound, realistic, and interpretable has motivated decades of research with a vast array of applications because satisfying all three of these goals simultaneously is challenging. Our proposed PCC framework represents a novel perspective in the construction of covariance functions. Through an application involving the movement of killer whales, we demonstrated that the PCC framework can be used to create sophisticated covariance functions that account for several important mechanisms, without resorting to unrealistic assumptions such as separability.

The joint movement of interacting individuals can be viewed as a multivariate temporal process. Methods for interpolating multivariate spatial and spatio-temporal processes, also called “cokriging,” have been studied for several decades (e.g., Myers, 1982; Cressie, 1991). The primary challenge has been to develop models for multivariate processes that accurately capture the complex forms of data within each process, as well as across processes, while ensuring the resulting cross-covariance structure is valid (symmetric and non-negative definite). While not explicitly framed as a PCC, Ver Hoef and Barry (1998) proposed a model for cokriging that shares important connections with our proposed model for the joint movement of interacting animals. Ver Hoef and Barry (1998) also approach the problem of specifying a valid, mechanistically motivated cross-covariance structure for multivariate processes through the use of kernel convolutions, allowing the authors to relax existing unrealistic assumptions about the covariance structure. Similar to the way dependence among individuals arises through the application of a social smoothing kernel (Section 2.3.1), Ver Hoef and Barry (1998) induced dependence in a multivariate process by smoothing across variables. The cross-covariance in Ver Hoef and Barry (1998) also accommodates spatio-temporal lags in the dependence among variables. This generalization may also be a useful feature in future models for animal movement, where a temporal lag may allow researchers to capture the effect of animals following one another.

The particular model we constructed for animal movement cannot be used to understand all forms of dependence that may exist among interacting individuals. Rather, it has been tailored to the case in which interactions among animals manifest themselves as movement along proximally close trajectories. Particles may also exhibit other forms of dependence, such as a tendency to repel in the case of strong territorial behavior, that would not be well-described by the same covariance function. The PCC framework allows researchers to define and order kernels as necessary to appropriately model the mechanisms under study. Under a different ordering, the same three kernels we employed result in another useful covariance function for the study of particle movement. We briefly discuss this alternative to highlight the flexibility offered by the PCC framework and provide some intuition for the role certain kernel functions play in the characteristics of the resulting random process.

As mentioned in Section 2.1, choosing the ordering of the kernels in a PCC represents a meaningful modeling decision when kernel convolutions do not commute. In our proposed model for dependent movement, the second and third stages of smoothing commute, while the first and second

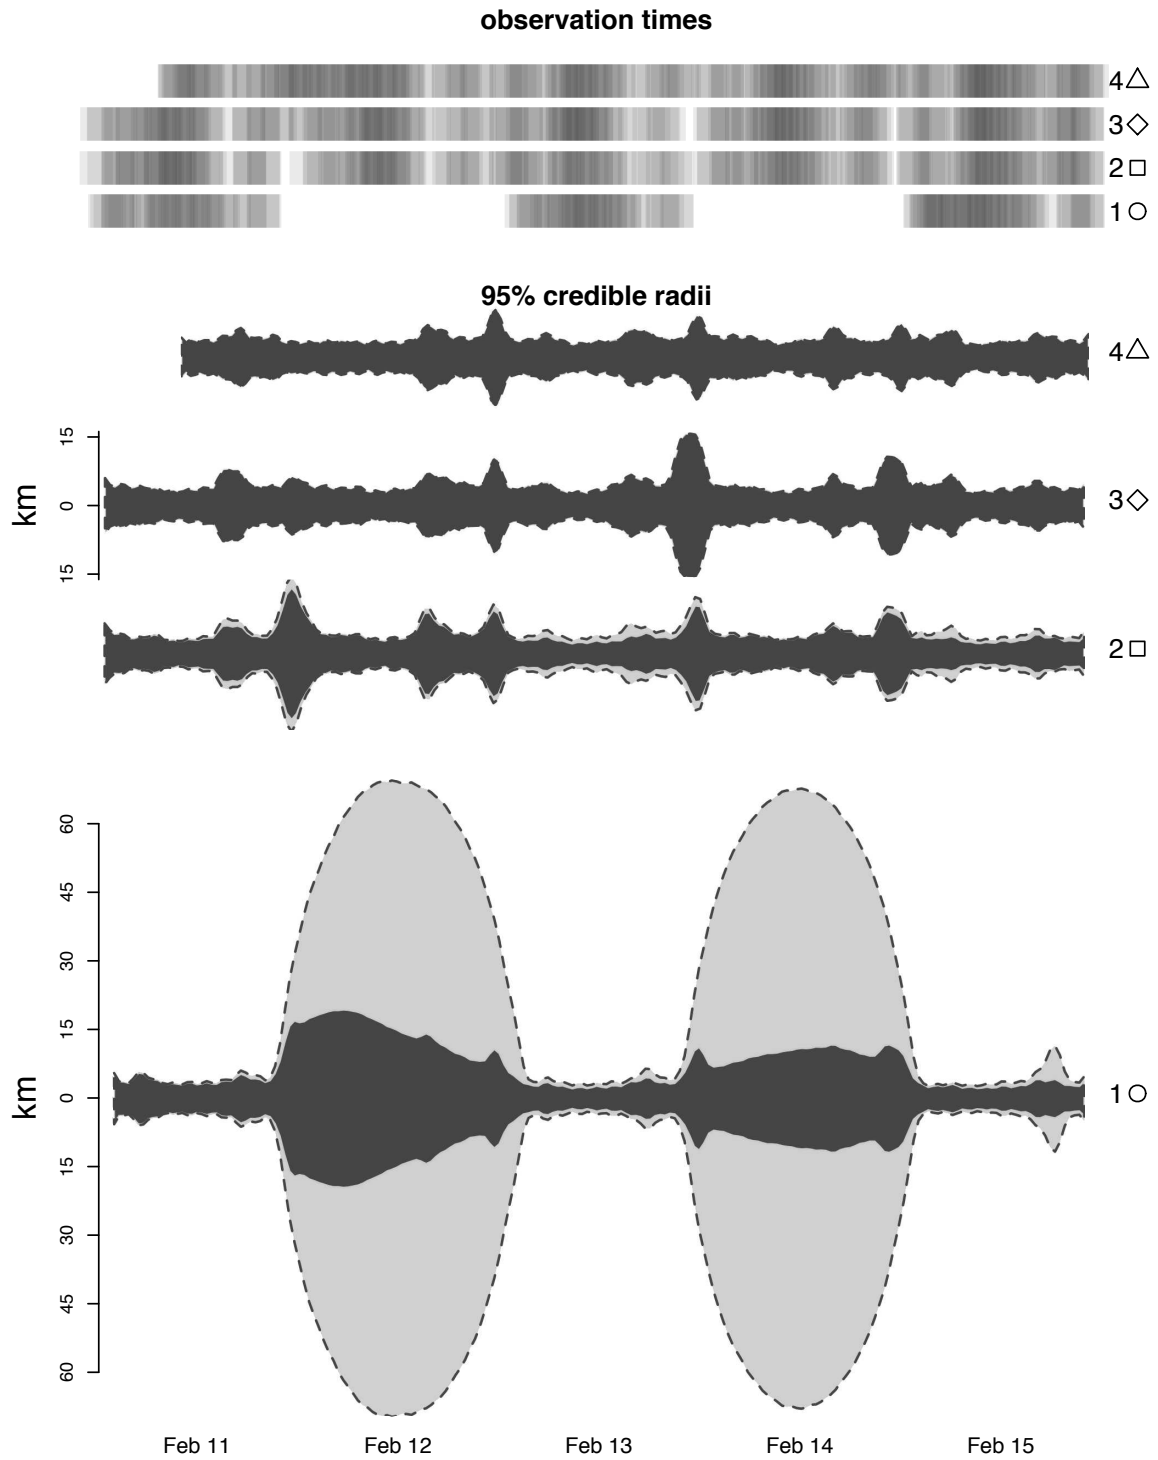


Fig. 4. *TOP:* Observation times for the three killer whales. Darker regions correspond to denser observation times, or equivalently, shorter gaps between observations. *BOTTOM:* 95% circular credible region radii at each time point for all four individuals. The dashed outer line shows the radii for the independent model, and the solid polygons show the radii for the full model.




Fig. 5. Posterior distribution of the dynamic social network that defines the social smoother used in stage two of the proposed PCC model. The dashed line represents the posterior median, and the polygon represents a pointwise 95% credible interval. Two of the six total connections are shown above; see Appendix 8.2.2 for others.

do not. We can interpret the commutativity between the social and inertial smoothers mechanistically, by noting that it does not matter whether the effect of the latent social connections operates on the raw, Brownian paths, or the temporally smoothed paths. In contrast, if we socially smoothed white noise and then used the kernel associated with Brownian motion, a very different form of dependence results. This alternative ordering ($h^{(soc)} \rightarrow h^{(bm)} \rightarrow h^{(inl)}$) presents another plausible mechanism through which dependence might arise in paths taken by multiple particles.

When applied to the velocity of a particle, smoothing with the step function kernel $h^{(bm)}$ returns the associated position process (after taking into consideration the appropriate initial location of the particle). Therefore, Brownian motion can be thought of as a random position process in which the velocity of the particle during each infinitesimal span of time is a realization from a Gaussian white noise process. In our proposed PCC model for dependent movement, the social smoother has the effect of “shrinking” the positions of strongly-connected particles toward each other. If we instead employed the social smoother before the step function kernel, the result would be to “shrink” the velocities of the particles together, rather than their positions. Thus, we would be inducing a tendency for particles to move in similar directions, though not necessarily a tendency to be in similar locations. Such an effect would be visible in the particles as movement in parallel, perhaps with a considerable distance between connected particles.

A PCC approach to constructing covariance functions for GPs allows for broad flexibility in model development, offering researchers a highly customizable framework that can be used in a wide variety of applications. We demonstrated the value of this approach with an application to animal movement, however PCCs can be used to model a broad range of random processes. Rather than relying on parametric families of covariance functions, PCCs encourage the use of interpretable, intuitive, and problem-specific convolution kernels, allowing for direct incorporation of scientific knowledge.

7. Acknowledgements

Killer whale tagging was conducted under permit #14097 from the National Marine Fisheries Service and Antarctic Conservation Act permit #2009-013. Shipboard tagging operations were supported by Lindblad Expeditions and the National Geographic Society, and by an NSF rapid grant to Ari Friedlaender. Robert Pitman helped with tag deployments and identification of killer whale types in the field.

This material is based in part upon research supported by NSF DMS-1614392. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government. The findings and conclusions of the NOAA authors in the paper are their own and do not necessarily represent the views of the National Marine Fisheries Service, NOAA.

References

- Andrews, R. D., Pitman, R. L. and Ballance, L. T. (2008) Satellite tracking reveals distinct movement patterns for Type B and Type C killer whales in the southern Ross Sea, Antarctica. *Polar Biology*, **31**, 1461–1468.
- Barry, R. P. and Ver Hoef, J. M. (1996) Blackbox Kriging: Spatial Prediction without Specifying Variogram Models. *Journal of Agricultural, Biological, and Environmental Statistics*, **1**, 297–322.
- Bolin, D. and Lindgren, F. (2013) A comparison between Markov approximations and other methods for large spatial data sets. *Computational Statistics & Data Analysis*, **61**, 7–21.
- Brost, B. M., Hooten, M. B., Hanks, E. M. and Small, R. J. (2015) Animal movement constraints improve resource selection inference in the presence of telemetry error. *Ecology*, **96**, 2590–2597.
- Buderman, F. E., Hooten, M. B., Ivan, J. S. and Shenk, T. M. (2016) A functional model for characterizing long-distance movement behaviour. *Methods in Ecology and Evolution*, **7**, 264–273.
- Calder, C. A. (2007) Dynamic factor process convolution models for multivariate spacetime data with application to air quality assessment. *Environmental and Ecological Statistics*, **14**, 229–247.
- (2008) A dynamic process convolution approach to modeling ambient particulate matter concentrations. *Environmetrics*, **19**, 39–48.
- Codling, E. A. and Bode, N. W. (2014) Copycat dynamics in leaderless animal group navigation. *Mov. Ecol.*, **2**, 11.
- Cressie, N. (1991) *Statistics for Spatial Data*. New York, New York, USA: John Wiley & Sons, first edn.
- Cressie, N. and Pavlicová, M. (2002) Calibrated spatial moving average simulations. *Statistical Modelling*, **2**, 267–279.
- Durante, D., Scarpa, B. and Dunson, D. B. (2014) Locally adaptive factor processes for multivariate time series. *Journal of Machine Learning Research*, **15**, 1493–1522.
- Durban, J. W., Fearnbach, H., Burrows, D. G., Ylitalo, G. M. and Pitman, R. L. (2016) Morphological and ecological evidence for two sympatric forms of Type B killer whale around the Antarctic Peninsula. *Polar Biology*, 1–6.
- Durban, J. W. and Pitman, R. L. (2012) Antarctic killer whales make rapid, round-trip movements to subtropical waters: evidence for physiological maintenance migrations? *Biology Letters*, **8**, 274–277.
- Feynman, R. P., Leighton, R. B. and Sands, M. (1963) *The Feynman Lectures on Physics*. Reading, Massachusetts: Addison-Wesley.
- Finley, A. O., Banerjee, S. and Gelfand, A. E. (2013) spBayes for large univariate and multivariate point-referenced spatio-temporal data models. *arXiv preprint arXiv:1310.8192*.

- Fleming, C. H., Fagan, W. F., Mueller, T., Olson, K. A., Leimgruber, P. and Calabrese, J. M. (2015) Estimating where and how animals travel: An optimal framework for path reconstruction from autocorrelated tracking data. *Ecology*, 15–1607.1.
- Gelfand, A. E., Kim, H.-J., Sirmans, C. F. and Banerjee, S. (2003) Spatial Modeling With Spatially Varying Coefficient Processes. *Journal of the American Statistical Association*, **98**, 387–396.
- Goldenberg, A., Zheng, A. X., Fienberg, S. E. and Airolidi, E. M. (2010) A Survey of Statistical Network Models. *Found. Trends Mach. Learn.*, **2**, 129–233.
- Goldenberg, S. Z., de Silva, S., Rasmussen, H. B., Douglas-Hamilton, I. and Wittemyer, G. (2014) Controlling for behavioural state reveals social dynamics among male African elephants, *Loxodonta africana*. *Animal Behaviour*, **95**, 111–119.
- Hansell, S. (1984) Cooperative Groups , Weak Ties , and the Integration of Peer Friendships. *Social Psychology Quarterly*, **47**, 316–328.
- Haydon, D. T., Morales, J. M., Yott, A., Jenkins, D. A., Rosatte, R. and Fryxell, J. M. (2008) Socially informed random walks: incorporating group dynamics into models of population spread and growth. *Proceedings. Biological sciences / The Royal Society*, **275**, 1101–1109.
- Higdon, D. M. (1998) A process-convolution approach to modelling temperatures in the North Atlantic Ocean. *Environmental and Ecological Statistics*, **5**, 173–190.
- (2002) Space and space-time modeling using process convolutions. In *Quantitative methods for current environmental issues*, 37–56. Springer.
- Hoff, P. D. (2008) Modeling homophily and stochastic equivalence in symmetric relational data. In *Nips*, 657–664.
- Hoff, P. D., Raftery, A. E. and Handcock, M. S. (2002) Latent Space Approaches to Social Network Analysis. *Journal of the American Statistical Association*, **97**, 1090–1098.
- Hooten, M. B. and Johnson, D. S. (2016) Basis Function Models for Animal Movement. *Journal of the American Statistical Association*, **In Press**.
- Hooten, M. B., Johnson, D. S., McClintock, B. T. and Morales, J. M. (2017) *Animal Movement*. Boca Raton, Florida, USA: CRC Press Taylor & Francis Group.
- Johnson, D. S., London, J. M., Lea, M.-A. and Durban, J. W. (2008) Continuous-time correlated random walk model for animal telemetry data. *Ecology*, **89**, 1208–15.
- Kahle, D. and Wickham, H. (2013) ggmap: Spatial Visualization with ggplot2. *The R Journal*, **5**, 144–166.
- Katzfuss, M. (2016) A multi-resolution approximation for massive spatial datasets. *Journal of the American Statistical Association*.
- Langrock, R., Hopcraft, J. G. C., Blackwell, P. G., Goodall, V., King, R., Niu, M., Patterson, T. A., Pedersen, M. W., Skarin, A. and Schick, R. S. (2014) Modelling group dynamic animal movement. *Methods in Ecology and Evolution*, **5**, 190–199.
- McClintock, B. T., Johnson, D. S., Hooten, M. B., Ver Hoef, J. M. and Morales, J. M. (2014) When to be discrete: the importance of time formulation in understanding animal movement. *Movement Ecology*, **2**, 21.

- Myers, D. E. (1982) Matrix formulation of co-kriging. *Journal of the International Association for Mathematical Geology*, **14**, 249–257.
- Nychka, D. W., Bandyopadhyay, S., Hammerling, D., Lindgren, F. and Sain, S. (2015) A Multiresolution Gaussian Process Model for the Analysis of Large Spatial Datasets. *Journal of Computational and Graphical Statistics*, **24**, 579–599.
- Paciorek, C. J. (2003) *Nonstationary Gaussian processes for regression and spatial modelling*. Ph.D. thesis, Carnegie Mellon University.
- Potts, J. R., Bastille-Rousseau, G., Murray, D. L., Schaefer, J. A. and Lewis, M. A. (2014) Predicting local and non-local effects of resources on animal space use using a mechanistic step selection model. *Methods in Ecology and Evolution*, **5**, 253–262.
- Rue, H. and Held, L. (2005) Continuous time random walks. In *Gaussian Markov Random Fields: Theory and Application*, chap. 3.5, 123–130. Chapman and Hall/CRC.
- Russell, J. C., Hanks, E. M. and Haran, M. (2016) Dynamic Models of Animal Movement with Spatial Point Process Interactions. *Journal of Agricultural, Biological, and Environmental Statistics*, **21**, 22–40.
- Scharf, H. R., Hooten, M. B., Fosdick, B. K., Johnson, D. S., London, J. M. and Durban, J. W. (2016) Dynamic social networks based on movement. *Annals of Applied Statistics*, **10**, 2182–2202.
- Shepp, L. A. (1966) Radon-Nikodym Derivatives of Gaussian Measures. *Annals of Mathematical Statistics*, **37**, 321–354.
- Ver Hoef, J. M. and Barry, R. P. (1998) Constructing and fitting models for cokriging and multi-variable spatial prediction. *Journal of Statistical Planning and Inference*, **69**, 275–294.
- Ver Hoef, J. M., Peterson, E. and Theobald, D. (2006) Spatial statistical models that use flow and stream distance. *Environmental and Ecological Statistics*, **13**, 449–464.
- Ver Hoef, J. M. and Peterson, E. E. (2010) A Moving Average Approach for Spatial Statistical Models of Stream Networks. *Journal of the American Statistical Association*, **105**, 6–18.
- Wecker, W. P. and Ansley, C. F. (1983) The signal extraction approach to nonlinear regression and spline smoothing. *J. Am. Statist. Assoc.*, **78**, 81–89.
- Wikle, C. K. (2010) Low-rank representations for spatial processes. In *Handbook of Spatial Statistics*, chap. 8, 107–118. Boca Raton, FL: CRC Press Taylor & Francis Group.
- Williams, R. and Lusseau, D. (2006) A killer whale social network is vulnerable to targeted removals. *Biology Letters*, **2**, 497–500.

8. Appendix

8.1. Model specification

Data model

$$\begin{aligned}
\mathbf{s} &\sim \mathcal{N}\left(\mathbf{0}, \sigma_s^2 \left(\mathbf{I} + \sigma_{\mu/s}^2 \Delta \tau \tilde{\mathbf{H}}' \Sigma_{dB} \tilde{\mathbf{H}}\right)\right) \\
\tilde{\mathbf{H}} &= \mathbf{H}^{(inl)}(\phi_{inl}) \mathbf{H}^{(soc)}(d\mathbf{B}_w, \phi_w, \sigma_w^2) \mathbf{H}^{(bm)} \\
\Sigma_{dB}(t, \tau, i, j) &= \mathbf{1}_{\{t=\tau\}} (\sigma_0^2 \mathbf{1}_{\{t=0\}} + \mathbf{1}_{\{t>0\}}) \mathbf{1}_{\{i=j\}} \\
\mathbf{H}^{(bm)}(t, \tau, i, j) &= \mathbf{1}_{\{\tau < t\}} \mathbf{1}_{\{i=j\}} \\
\mathbf{H}^{(soc)}(t, \tau, i, j) &= \mathbf{1}_{\{\tau=t\}} \frac{w_{ij}(t)}{|w_i(t)|} \\
\mathbf{H}^{(inl)}(t, \tau, i, j) &= \exp\left\{\frac{-(t-\tau)^2}{2\phi_{inl}^2}\right\} \mathbf{1}_{\{i=j\}}
\end{aligned}$$

Process model (integrated out for model fitting)

$$\begin{aligned}
\boldsymbol{\tau} &\equiv (\tau_1 = 0, \tau_2, \dots, \tau_{m-1}, \tau_m = 1)' \\
\Delta\tau_i &\equiv \tau_i - \tau_{i-1}, \quad 1 < i \leq m \\
dB(\tau_i) &\sim \begin{cases} \mathcal{N}(0, \sigma_0^2); & i = 1 \\ \mathcal{N}(0, \Delta\tau_i); & i > 1 \end{cases} \\
d\mathbf{B} &= (dB(\tau_1), \dots, dB(\tau_m))' \\
\tilde{\boldsymbol{\mu}} &= \mathbf{H}^{(inl)} \mathbf{H}^{(soc)} \mathbf{H}^{(bm)} d\mathbf{B}
\end{aligned}$$

Prior model

$$\begin{aligned}
\phi_{inl} &\sim \text{Gamma}(\alpha_s, \beta_s) & \sigma_0^2 &\sim \text{IG}(a_0, b_0) \\
\sigma_{\mu/s} &\sim \text{IG}(a_{\mu/s}, b_{\mu/s}) & \sigma_s^2 &\sim \text{IG}(a_s, b_s)
\end{aligned}$$

Fixed parameters

$$\phi_w \text{ and } \sigma_w^2 \tag{27}$$

8.2. MCMC details

8.2.1. Simulation Study

We simulated approximately continuous true paths using a grid of 500 equally spaced time points on the unit interval. We simulated 100 observation times drawn uniformly from the same unit interval. We set the hyperparameters associated with the network to $\phi_w = 4/15$ and $\sigma_w^2 = 10$. For each fit, we obtain 2,000 iterations and discard the first 1,000 as burnin.

8.2.2. Killer whales

We acquire 100,000 iterations on a single computing node, and discarded the first 75,000 as burnin to yield a sample size of 25,000. We employ diffuse priors for all parameters. We set the hyperparameters associated with the network to $\phi_w = 0.3$ and $\sigma_w^2 = 10$.

Table 1. True values and prior distributions used in simulation study.

parameter	true	prior density
ϕ_{inl} (“low tortuosity”)	0.04	Gamma(2, 100)
ϕ_{inl} (“high tortuosity”)	0.04/3	Gamma(2, 100)
σ_0^2	1	IG(10^{-3} , 10^{-3})
$\sigma_{\mu/s}^2$	800	IG(10^{-3} , 10^{-3})
σ_{μ}^2	10	
σ_s^2	0.0125	IG(10^{-3} , 10^{-3})


**Fig. 6.** Posterior distribution of the dynamic social network that defines the social smoother used in stage two of the proposed PCC model (IP-DEP). Each plot shows the time-evolution for a specific pair of killer whales indexed by row and column.

Table 2. Posterior credible intervals and prior distributions for the IP-IND model.

parameter	posterior		prior
	median	(2.5%, 97.5%)	density
ϕ_{inl}	0.0102	(0.0095, 0.0109)	Gamma(2, 100)
σ_0^2	1990	(910, 5900)	IG(1, 10)
$\sigma_{\mu/s}^2$	1320	(1060, 1660)	IG(10^{-3} , 10^{-3})
σ_{μ}^2	1.70	(1.39, 2.10)	NA
σ_s^2	0.00129	(0.00119, 0.00141)	IG(10^{-3} , 10^{-3})

Table 3. Posterior credible intervals and prior distributions for the IP-DEP model.

parameter	posterior		prior
	median	(2.5%, 97.5%)	density
ϕ_{inl}	0.0102	(0.0095, 0.0109)	Gamma(2, 100)
σ_0^2	1450	(630, 4170)	IG(1, 10)
$\sigma_{\mu/s}^2$	1850	(1430, 2420)	IG(10^{-3} , 10^{-3})
σ_{μ}^2	2.38	(1.87, 3.06)	NA
σ_s^2	0.00128	(0.00118, 0.00139)	IG(10^{-3} , 10^{-3})